

1/13

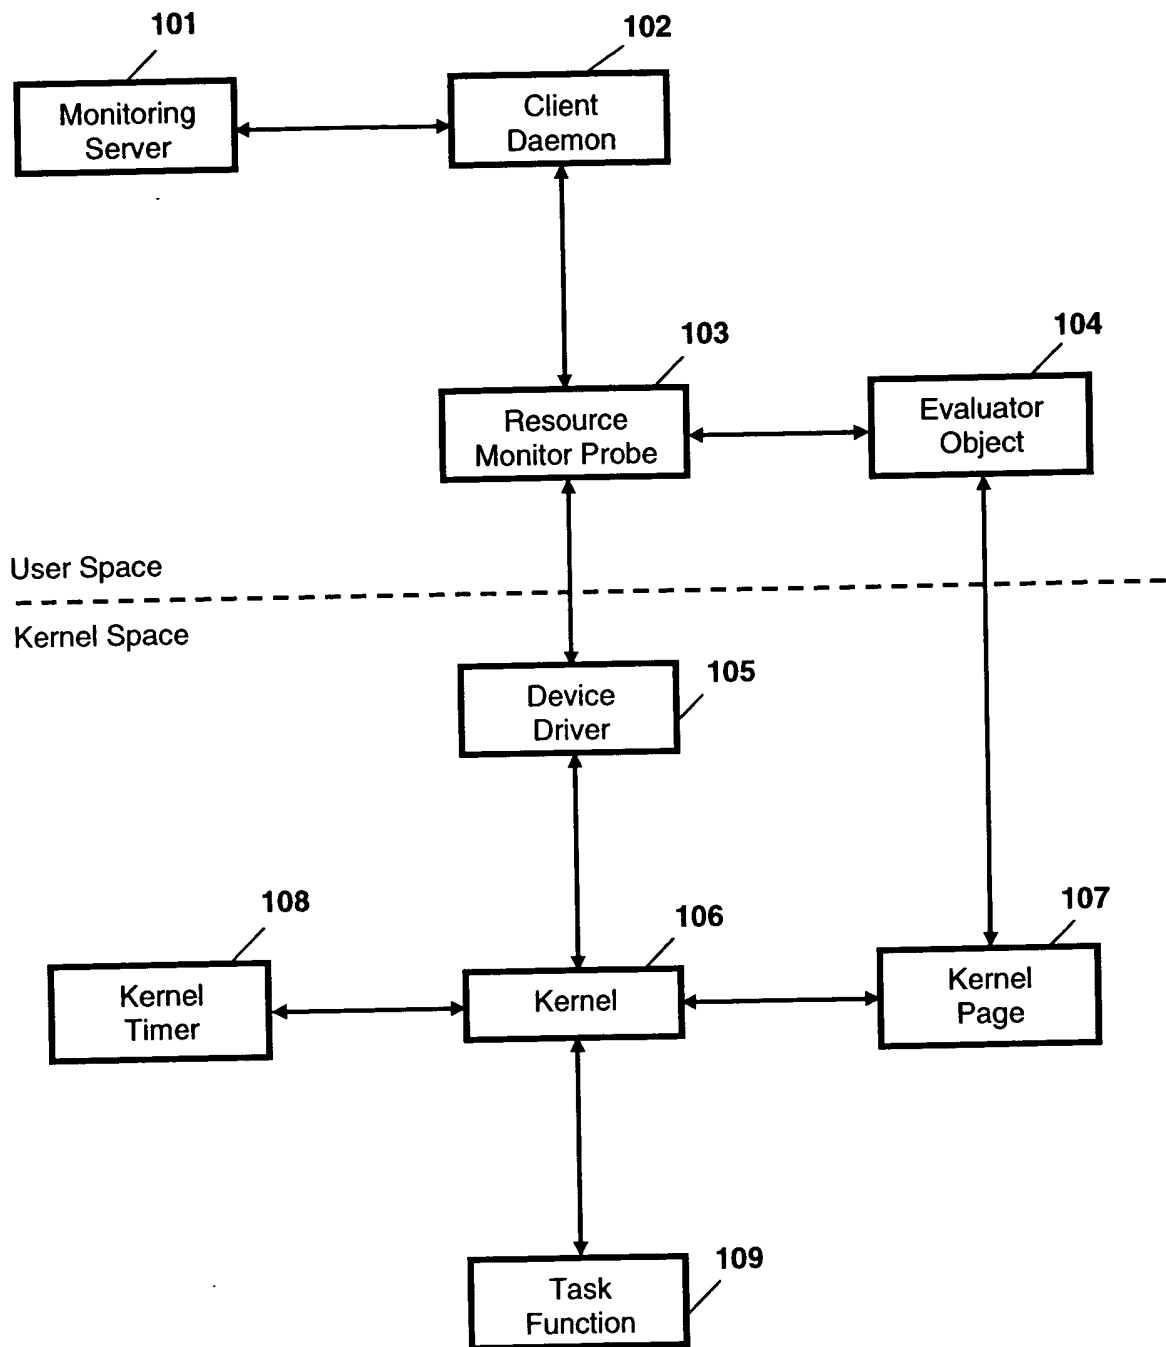


FIG. 1

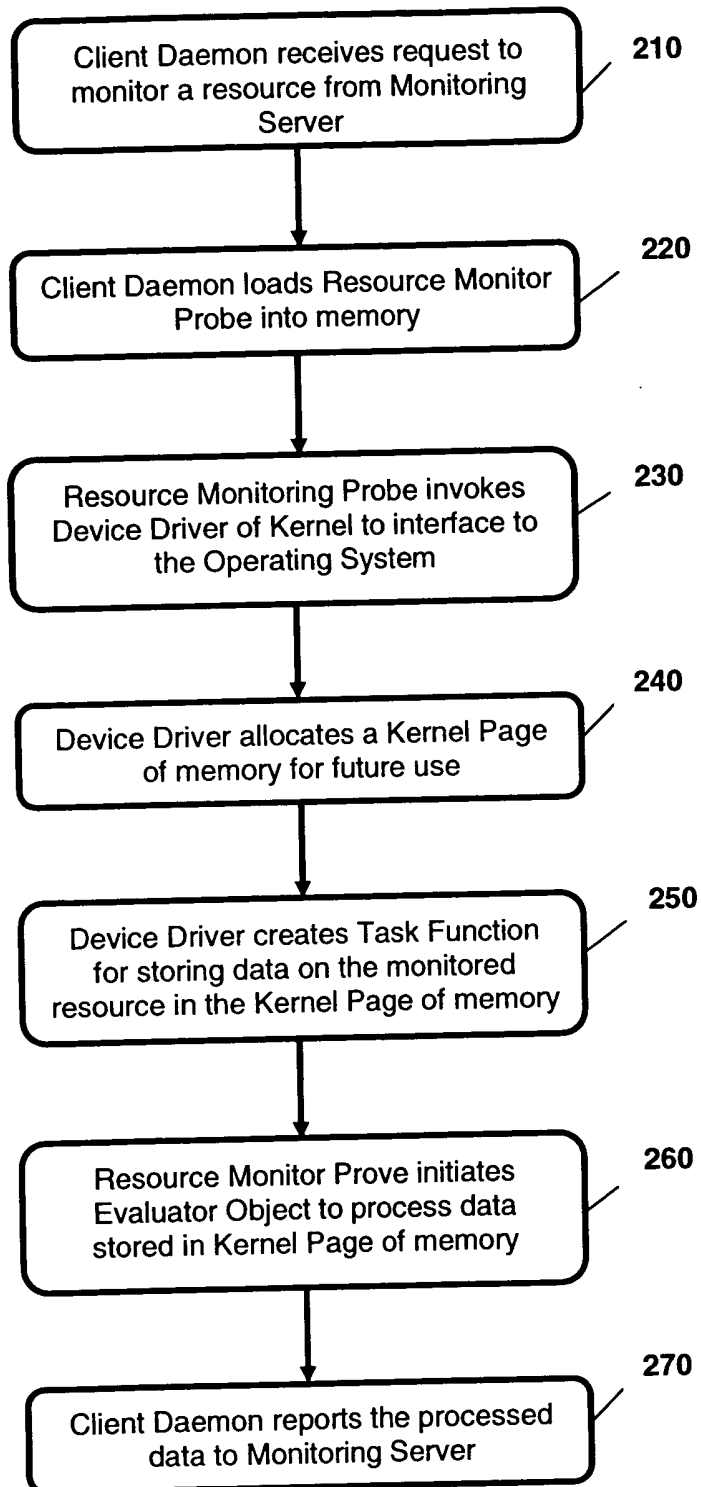


FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7

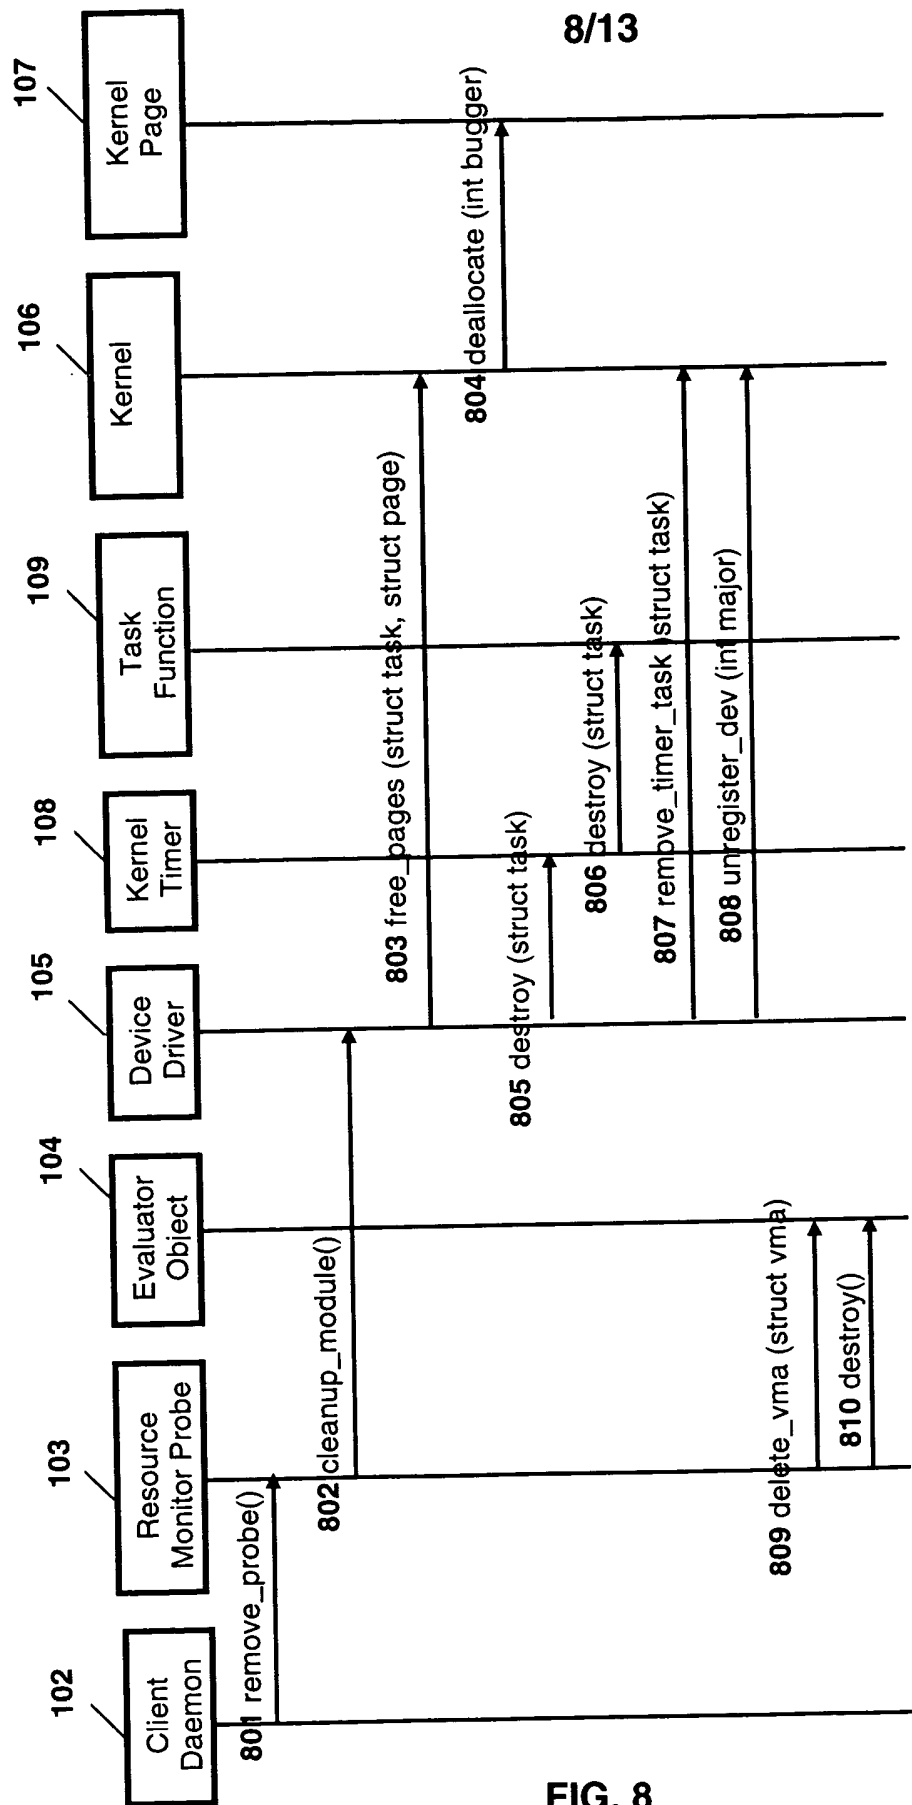


FIG. 8



```

#include <linux/module.h>
#include <linux/fs.h>
#include <linux/mm.h>
#include <linux/malloc.h>
#include <linux/sched.h>

#define PROBE_MAJOR 0

struct file_operations cpu_probe_fops = {
    read: probe_read,
    write: probe_write,
    open: probe_open,
    release: probe_release,
    mmap: probe_mmap,
};

void *buffer;    // place to store the physical address of the kernel pages
void *head_buffer; //current location to be written by our timer function
struct semaphore sem;

//use dynamic assignment
int probe_major =PROBE_MAJOR;
int order=0; //request for one page only

int init_module(void)
{
    result=register_chrdev(probe_major,"cpu_probe",&cpu_probe_fops);
    if(result<0) {
        printk(KERN_WARNING"cpu_probe: can't get major %d
\n",probe_major);
        return results;
    }
    if (probe_major==0) probe_major=result;

    buffer = (void*)__get_free_pages(GFP_KERNEL,order)

```

FIG. 9A

```

        if(!Buffer)
            goto nomem;

        memset(buffer,0,PAGE_SIZE<<order);
        return 0;

nomem:
        unregister_chrdev(probe_major,"cpu_probe");
        return 1;
    }

void cleanup_module(void)
{
    if(buffer)
        free_pages((unsigned long)(buffer),order);

    unregister_chrdev(probe_major,"cpu_probe");
}

int probe_open(struct inode * inode, struct file *filp)
{
    MOD_INC_USE_COUNT;
    return 0;
}

int probe_release(struct inode *inode, struct file *filp)
{
    MOD_DEC_USE_COUNT;
    return 0;
}

int probe_read(struct file *filp, char * buf, size_t count, loff_t *fpos)

```

FIG. 9B

```

{
    if(copy_to_user(buf,(char*)buffer,PAGE_SIZE)))
        return -EFAULT;
    return PAGE_SIZE;
}

int probe_mmap(struct file *filp,struct vm_area_struct *vma)
{
    unsigned long offset=buffer;

    if(offset>=_&thinsp;_pa(high_memory))||(filp->f_flags&O_SYNC))
        vma->vm_flags |=VM_IO;
    vma->vm_flags |= VM_RESERVED;
    If(remap_page_range(vma->vm_start,offset,vma->vm_end - vma->vm_start,
        return -EAGAIN;
return 0;
}

// pages[order] array of number of pages
/* the memory address returned by kmalloc and get_free_pages are also
virtual addresses.their actual value is still massaged by the MMU before
it is used to address physical memory */

wait_queue_head_t probe_wait;
struct timer_list probe_timer;

void probe_read_pid(unsigned long ptr)
{
    written =sprintf((char*)head_buffer,"%08u",(int)(current->pid));
    probe_data.len+=sizeof(int);
    probe_data.buf=buffer+sizeof(int);
    wake_up_interruptible(&probe_wait);
}

```

FIG. 9C

```
void probe_write(struct file *filp, char *buf, size_t count, loff_t *fpos)
{
    probe_data.len=0;
    probe_data.buf=buffer;
    probe_data.jiffies=jiffies;
    probe_data.queue=NULL;

    init_timer(&probe_timer);
    probe_timer.function = probe_read_pid;
    probe_timer.data=(unsigned long) &probe_data;
    probe_timer.expires=jiffies+HZ;    //one sec timeout
    add_timer(&probe_timer);
    interruptible_sleep_on(&probe_wait);
    del_timer_sync(&probe_timer);

    return 0;
}
```

FIG. 9D

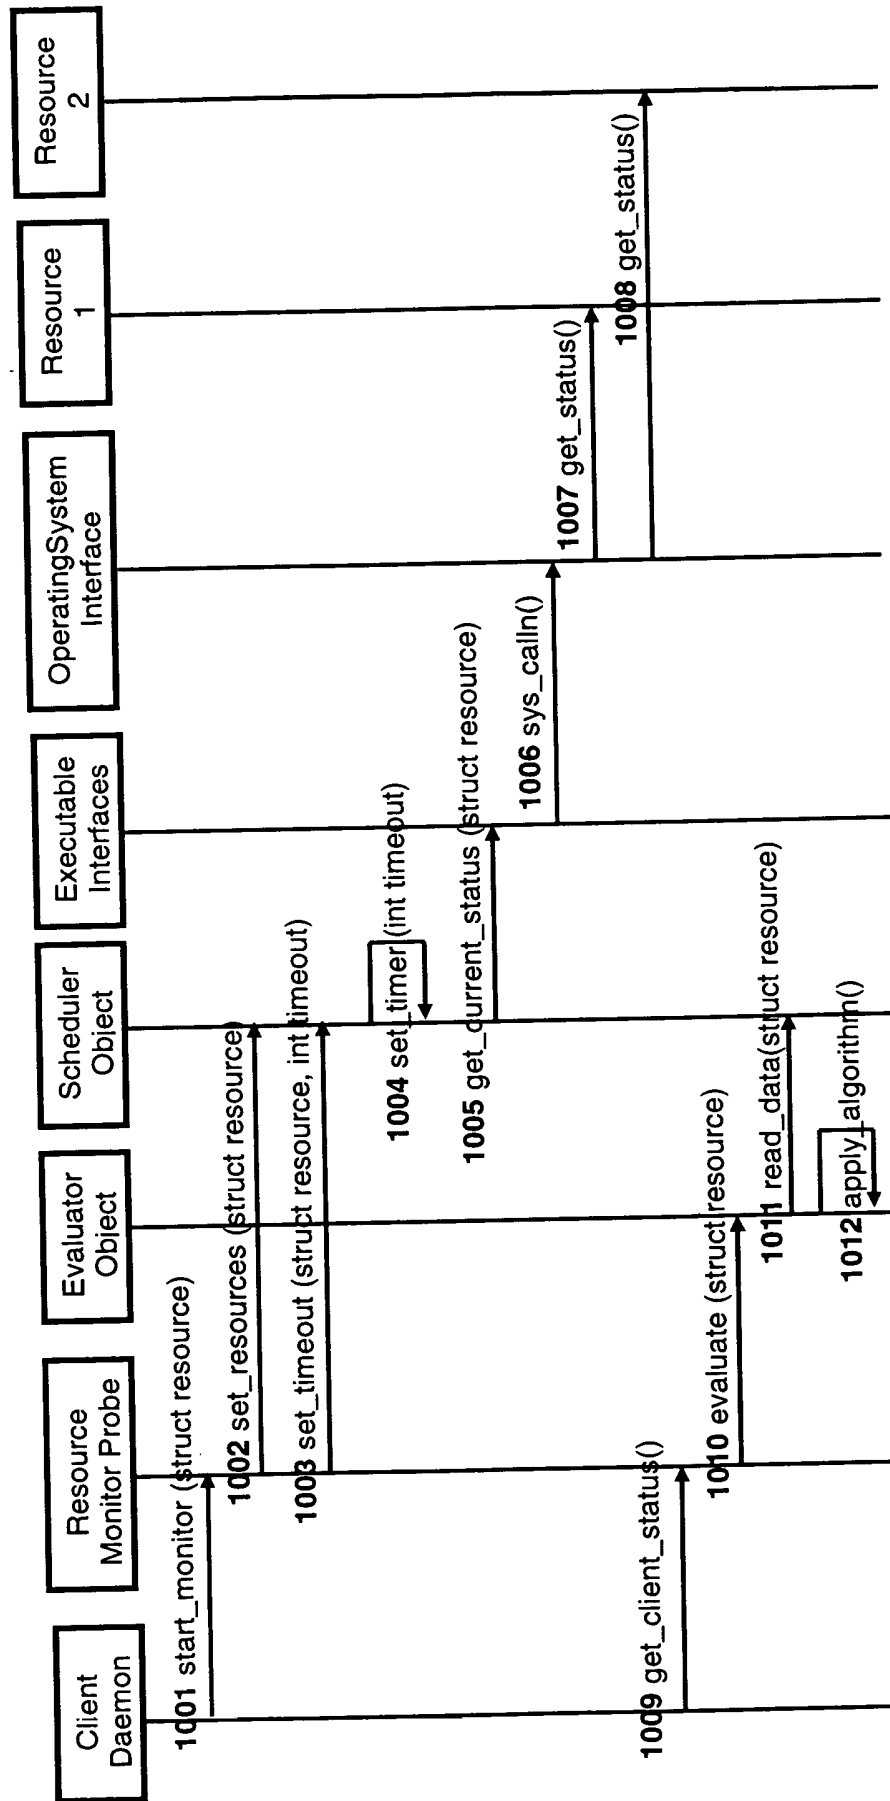


FIG. 10